

The Malicious Bot Playbook: Early Warning Signs, and What to Do About Them

Introduction

Bots account for nearly 40 percent of online traffic today – and many of those bots are out to damage organizations like yours. From content scraping to inventory hoarding to credential-stuffing, malicious bot actors are growing more complex and sophisticated every year. Many are even able to bypass CAPTCHAs and other simple user verification tests.

What's more, while some bots sharpen their tactics to target specific types of businesses, a single bot or botnet can threaten many different types of organizations.

For this reason, no single tactic can stop every bot, and prevent it from harming your users and your brand.

The only effective approach is to stay alert for a diverse range of telltale bot warning signs – and respond to each by gathering data and then deploying targeted responses, pattern detection, predictive analytics, and other complementary strategies.

Warning Signs of a Bot Problem

By tracking a range of potential indicators, you stand an excellent chance of spotting bad bots before they have a chance to do serious damage. Here is what to look for:

Higher infrastructure costs with no increase in business

All traffic to your website carries some cost. No matter who or what accesses your content, you've got to foot the bill for storage and compute. But bad bots can increase your traffic-related costs without providing any revenue to your business. While good bots are used by search engines to index content on your site, and thus support your SEO rank, bad bots run up significant excessive bandwidth charges every year.

Unusual purchases of low-volume, high-demand inventory

If you notice that you're selling a suspiciously high percentage of your inventory to a surprisingly small subset of buyers, inventory-hoarding bots may be the culprits. While some of these bots will simply fill and abandon shopping carts in order to block legitimate customers, others will actually buy your inventory with the goal of reselling it for a higher price on other sites.

Increased customer complaints

An uptick in support tickets related to account lockouts and fraudulent transactions could be a sign of credential-stuffing bots. These bots take over legitimate user accounts with information they've harvested from past leaks. In addition to negatively impacting your customer experience, these fraudulent transactions will overload your servers — creating longer page load times — or even render your website unavailable.

Increase in failed login attempts

Every customer mistypes their password now and then — but if you see a sudden rash of failed login attempts, you've very likely got a bot problem. While some credential-stuffing bots try to access legitimate customer accounts via stolen credentials, a simpler and more common technique is to launch a brute-force attack, in which bots attempt many rapid-fire logins using dictionaries of thousands of popular usernames and passwords. When a bot exceeds your site's limit of failed logins for a particular account, that account's real human owner will be locked out until you resolve the issue — a major user experience headache.



Low yield on advertising spend

Digital advertising can be an effective tool for driving traffic to your site, but it's also a lucrative weapon for bad bots. Many traffic bots mimic the behavior of human users – clicking your ads repeatedly to drive up your pay-per-click (PPC) spend, then bouncing without making a purchase. While some advertising platforms have deployed machine-learning algorithms to cut down on click fraud, much of it remains undetected¹. That's why it's crucial to get proactive, and monitor every click that comes through your ads.



Skewed page-view analytics

If your page views suddenly spike upward for no discernable reason, bad bots could be the culprit. While a spike in traffic may come from human users if you've just launched a new product or an event promotion, malicious bot operators are getting smarter about deploying content-scraping bots at exactly these times, stealing your content and negatively affecting your aggregate analytics data.

Sudden increase in account creation

When hundreds or even thousands of new user accounts appear out of the blue, bots may be behind the influx. They can use these fake profiles to spam your public ratings and commit numerous other forms of fraud – threatening not only your revenue and user retention, but also your brand's credibility.

Duplicates of your content on non-approved sites

Other sites sharing your content can be good, but a sudden increase in outright duplicated content is a hallmark of content-scraping bots. These bots steal information you've taken the time to assemble and curate, allowing malicious site operators to host it on domains they own – boosting their own traffic while yours takes a hit.

¹ <https://www.entrepreneur.com/article/313943>

Traffic originating from unusual geographic locations

Sudden spikes originating from unexpected locations may point to bad bot activity – especially if this activity appears in clusters, centered in regions where your customers don't live or your services aren't available. Keep a close eye on any suspicious activity that appears unrelated to your regular user base.

Traffic from typical locations at unusual times

Just as traffic spikes from unexpected locations may point to malicious bots, spikes from normal locations at unusual times may indicate bots trying to disguise themselves as your regular users. If you see a spike in activity from a common region in the middle of the night, for example, you may want to investigate that traffic more closely.

Increase in card validation failures

A particularly dangerous sign of bad bots is an uptick in credit card transactions that fail to validate. Credit card stuffing bots will test thousands of stolen credit card numbers in an attempt to find one that works. They do this by making low-value purchases on less-secure websites, before performing larger transactions on bigger sites, or selling the validated card numbers on the dark web. Your site may factor into these plans at any point in the chain – and if the failed transactions are egregious enough, your payment provider may fine you.

Tactics for Fighting Bots

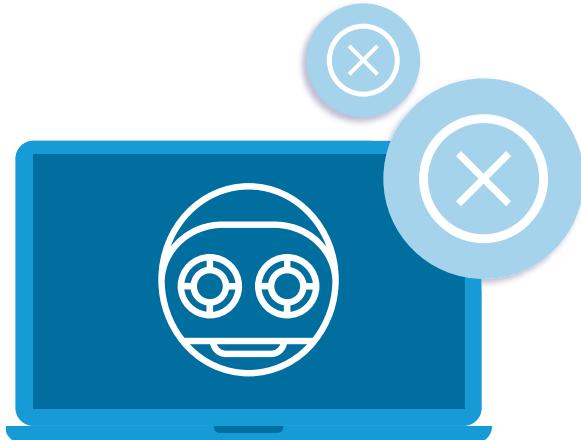
Just as no two bot attacks are alike, you'll usually need a combination of multiple tactics to stop them all in their tracks. Consider some of the following strategies:

Block bad bots as soon as you catch them

The most self-evident response to a bot is also one of the most effective: simply block all traffic that you've identified as coming from malicious bot activity.

This tactic alone can save you significant costs on bandwidth and storage – not to mention safeguarding consumer trust as well as your brand's reputation.

At the same time, remember that if a bot operator is highly motivated, they might change their tactics and come back later with a different attack strategy.



Whitelist all the good bots you're aware of

Even as you detect and block bad bots, it's crucial to make sure good bots from search engines and partners are still able to scrape your site. This not only ensures that your SEO ranking remains solid, but also keeps legitimate customer traffic flowing smoothly from third-party services that refer traffic to you. Whitelisting also makes it much easier to set bot-blocking rules that won't negatively impact access for real human visitors.

Challenge suspected bots you detect

As soon as you notice a pattern of suspicious logins, low page depth or time-on-page, failed credit card validations, or any other trademark bot behavior, it's crucial to deploy a security test such as CAPTCHA to verify that suspicion. However, note that many advanced bots can now solve these puzzles just as easily as human users. This means that while CAPTCHAs can be useful against some bots, they need to be combined with other tactics. What's more, CAPTCHAs create friction in your user experience, negatively impacting human visitors' perception of your brand. Thus, it's essential to ensure that only suspected bots – not legitimate users – are served CAPTCHAs.

Limit the rate at which users can request information

Rate limiting can be an effective technique for keeping less sophisticated bots at bay. By setting hard limits on the number of times any IP address can submit requests to your site, you'll prevent many simplistic brute-force bot attacks, which try to sign in using thousands of dictionary words and common passwords. However today's more advanced bots can keep their number of requests just below your rate limit— remaining undetected while they continue to inflict damage.



Keep detailed logs of all site traffic

While you're likely already maintaining daily logs of page views and account logins, more detailed logs of user information – such as IP addresses, browsers, devices, OSes, geolocations, referrers, networks and pageviews – can prove invaluable for detecting more subtle activity patterns. Logs can give you a clear idea of how bots tend to behave on your site, enabling you to set up more effective security policies. In addition, logs are often essential for reporting and compliance in the event that you do actually suffer a data breach.

Redirect bots to alternative content

When you're fairly certain a certain traffic source is a bot, serve it an alternate piece of content that consumes its computational resources. You can even feed fake data – such as erroneous pricing information – to content-scraping bots, rendering them useless to their operators. Techniques like these will buy you time to watch how each bot behaves, understand its activity pattern, and prepare a strategy to deal with it once and for all.

Require additional authentication for all users

As bot attacks become more prevalent, a growing number of sites are turning to heightened security measures, even for legitimate human logins. Two-factor authentication (2FA), for example, requires users to confirm their identity on multiple devices or accounts, while one-time passwords (OTPs) can discourage credential-stuffing bots by making accounts tougher to crack. But remember that these techniques may negatively impact your user experience by adding friction to your login process.



Conclusion

As you explore these tactics, remember: none of them will stop all bots as a standalone approach. To protect your business, you'll likely need to combine tactics — and to add advanced multivariable pattern analysis.

Cloudflare Bot Management can help organizations across a variety of industries adopt this multi-pronged approach. It's incorporated into the broader Cloudflare network, which supports over 25 million Internet properties and spans more than 200 global cities. By drawing on continuous threat intelligence from across that network, Cloudflare Bot Management offers behavior analysis, machine learning, and fingerprinting to remove much of the effort of combating bad bots. Visit www.cloudflare.com to learn more.

1 888 99 FLARE | enterprise@cloudflare.com | www.cloudflare.com

© 2020 Cloudflare Inc. All rights reserved.
The Cloudflare logo is a trademark of Cloudflare. All other company and product names may be trademarks of the respective companies with which they are associated.